parallel buses, etc. The I^2C signaling interface is a known signaling interface and comprises a clock line and a bi-directional data line. The use of such a known signaling interface provides an open, common, interconnect to each of the *N* boards. It should be noted that in this embodiment system maintenance bus 7 is separate from other signaling buses of server 5. As further described below, this provides for fault tolerant operation but is not essential to the inventive concept, i.e., the signaling represented by system maintenance bus 7 could be a part of another signaling bus of server 5.

Boards 200-1, 200-2, …, 200-N are illustratively represented by board 200, which is shown in FIG. 2. Like FIG. 1, other than the inventive concept, the elements shown in FIG. 2 are well known and will not be described in detail (e.g., the controller is a PIC16LC774 available from Microchip). Board 200 comprises a power control element 100 and a remainder of the board, as represented by central processing unit (CPU) 250 (hereinafter the "remainder of the board" is simply referred to as "the board"). Power control element 100 comprises at least one DC (direct current)-to-DC voltage regulator 110 and a controller 120. The latter comprises a stored-program controlled processor as represented by CPU 180 and memory 185. Voltage regulator 110 receives the input voltage(s) via power path 11 and provides at least one regulated voltage to the board (e.g., CPU 250), via signal path 112. In other words, without the power provided via signal path 112 – the board will not function. In contrast, power path 12 provides the keep alive voltage(s) to power control element 100 – thus, power control element 100 has power to function even if no power is provided, via power path 11, to the board. (It should be noted, that the keep alive power may be controlled elsewhere in server 5, and, in fact, may be individually controlled for each power control element.)

As described further below, controller 120 monitors power-related parameters of board 200 and adjusts the operation of voltage regulator 110 in response thereto. For the purpose of this example, the power-related parameters are represented by the signal level(s) of the regulated voltage(s) conveyed on signal path 112. However, the inventive concept also applies to other types of power-related parameters. For example, temperature – which is related to the amount of power being dissipated by the board – may be monitored either directly from voltage regulator 110 (if the voltage regulator

supports such a feature) and/or via a separate temperature sensing circuit (not shown). (Temperature sensing circuits are known in the art and not described herein.)

In the context of this example, controller 120 is coupled to signal path 112 for the purpose of monitoring the signal(s) conveyed thereon, and controls voltage regulator 110 via power path 121, which conveys at least one control signal. (For the purposes of this description, it is assumed that power path 121 is unidirectional. However, it should be noted that power path 121 could be bi-directional, e.g., additional status information may be provided from voltage regulator 110 to controller 120 for the purpose of conveying other power-related information such as, but not limited to, a current temperature level of voltage regulator 110, as mentioned above. Similarly, controller 120 may receive power-related information from other parts of board 200, e.g., a temperature sensor located at a position different from voltage regulator 110.)

Turning now to FIG. 3, an illustrative flow chart is shown in accordance with the principles of the invention. The inventive concept is implemented using conventional programming techniques, which as such, will not be described herein. In accordance with the flow chart of FIG. 3, controller 120 monitors power-related parameters of board 200 and, if necessary, controls voltage regulator 110 in response thereto. In particular, controller 120 executes steps 305 and 310 as long as a problem is not detected in step 310. However, once a problem is detected in step 310, controller 120 adjusts voltage regulator 110 in step 315 such as to ameliorate the detected problem.

For example, one power-related parameter is a regulated voltage level on signal path 112, and an illustrative problem is the regulated voltage level exceeding a predefined operating range. (The predefined operating range is either determined empirically and/or taken from data sheets for the various components used on board 200. For example, the predefined operating range may represent a guard band that is narrower than an operating range supported by the component data sheets.) As such, in the context of the flow chart of FIG. 3, if the regulated voltage level exceeds a predefined operating range (e.g., due to a variation in the effective impedance level, or load, presented by board 200) (steps 305 and 310), then controller 120 (in step 315) adjusts voltage regulator 110 in an attempt to keep the regulated voltage level within the predefined operating range. (Such variation

over time of the load presented by board 200 may occur due to, e.g., the particular type of operation being performed by board 200 at any point in time, as well as the aging of the board components (devices) themselves over time.)

A variation of the flow chart of FIG. 3 is shown in FIG. 4. Like numbers represent similar steps and are not described further. In the flow chart of FIG. 4, steps 350 and 355 have been added. In particular, controller 120 tracks (e.g., via the use of state variables) whether a previously detected problem still exists, i.e., has a "fault" occurred. As such, in step 350, controller 120 determines if board 200 has a fault (this could either be the same problem as detected previously, or just the fact that some problem still exists). Controller can detect a fault in any number of ways. For example, by tracking if a problem exists, or is repeating, beyond a predefined amount of time; or simply counting when a number of subsequent passes through step 350 exceeds a predetermined value, *K*. (The value of *K* is predefined and is assumed to be determined empirically or by other system parameters.) Once a fault is declared, controller 120 executes step 355 rather than step 315 (which attempts to clear the problem detected in step 310). In step 355, controller 120 illustratively shuts down board 200 since, apparently, a problem state cannot be cleared. Thus, a computer system supporting distributed power control has the ability to selectively shutdown individual boards. (As will be described below, other forms of exception handling can be performed, e.g., generating an alert, etc.)

As described above, and in accordance with the invention, each of the *N* boards has the ability to control the power distribution to itself – thus providing distributed power control for the server. Turning now to FIG. 5, another view of server 5 is shown illustrating another feature of the invention. Other than the inventive concept, the elements shown in FIG. 5 are well known and will not be described in detail. As noted above, and shown again in FIG. 5, each of the *N* boards, board 200-1 through 200-N, is coupled to system maintenance bus 7 (using the above-mentioned I^2C signaling). System maintenance bus 7 is coupled to port 510 for coupling to a maintenance, or administration, console 520, via path 501. The latter is representative of any coupling method, e.g., a wired and/or wireless connection, for coupling a terminal to a computer